

# A Comparison of MAC Protocols for Hybrid Fiber/Coax Networks: IEEE 802.14 vs. MCNS

*N. Golmie \*   F. Mouveaux   D. Su*

National Institute of Standards and Technology  
Gaithersburg, MD 20899

## Abstract

Both the IEEE and MCNS Medium Access Control (MAC) interface specifications for Hybrid Fiber/Coax networks share a common goal to provide a standard interface to the same anticipated services (such as cable telephony, Internet and on-line access for Web browsing, chat rooms, interactive games, local area network emulation for work-at-home applications and desktop conferencing). However the solutions offered by the two groups have little in common at the MAC layer. IEEE 802.14 chose ATM transfer as its default solution while MCNS uses a scheme that favors the delivery of variable length IP packets. Our goal in this study is to conduct a comparative performance evaluation of both IEEE 802.14 and MCNS MAC specifications. We focus on the efficiency of the two solutions in terms of contention access, ATM vs IP transfer, and adequate Quality of Service provision.

**Key Words:** Hybrid Fiber/Coax networks, Residential Broadband Access, Medium Access Control protocol.

## 1 Introduction

The IEEE 802.14 Cable TV Media Access Control (MAC) and Physical (PHY) Protocol Working Groups were formed in May 1994 by a number of vendors to develop international standards for data communications over cable. The original goal was to submit a cable modem MAC and PHY standards to the IEEE LAN/MAN Standard Committee in December 1995. However as of the writing of this paper, the IEEE 802.14 specifications remain a draft [3].

Tired of waiting for the IEEE 802.14, cable operators combined their purchasing power and operating under a limited partnership dubbed Multimedia Cable

Network System Partners Ltd. (MCNS) released the Data Over Cable System Interface Specification (DOCSIS) [4] for cable modem products in March 1997. The MCNS DOCSIS was first ratified by the Data Standard Subcommittee of the Society of Cable Telecommunications Engineers (SCTE) and then was approved as a new International Telecommunication Union (ITU) recommendation in March 1998. To date several vendors have announced plans to build products based on the MCNS DOCSIS standard.

At the physical layer which defines modulation formats for digital signals, the IEEE and the MCNS specifications are similar. The 802.14 specification supports ITU's J.83 Annex A, B and C standards for 64/256 QAM modulation on the downstream while MCNS supports Annex B only which is also the North American standard. Both 802.14 and MCNS support QPSK and 16 QAM for the upstream modulation. However the MAC layer specifications are substantially different. 802.14 chose ATM transfer as its default solution because it provides Quality of Service (QOS) required for integrated delivery of video, voice and data traffic. ATM is a long-term solution that has the flexibility to provide more than just Internet access. On the other hand, MCNS uses a scheme that favors the delivery of variable length IP packets rather than ATM in an attempt to keep cost and complexity of cable modems down.

Since cable products will soon be available supporting one or more standards (IEEE 802.14, MCNS or any other vendor specific solution) it becomes important to clearly identify the differences between the two emerging standards. Our motivation in this paper is to conduct a performance evaluation between the IEEE 802.14 and the MCNS MAC protocols. The questions we would like to answer are the following. Which MAC solution is best suited for HFC networks? Which of the MAC protocols best supports QOS in particular CBR end-to-end delay bounds and jitter? We compare the performance of both the IEEE 802.14 and the MCNS MAC specifications in terms of access delay, cell delay

---

\*Corresponding Author, Email: nada.golmie@nist.gov, Tel. (301) 975-4190, Fax: (301) 590-0932

variation, and probability density delay.

The rest of the paper is structured as follows. In section 2, we give some general overview of the HFC system and MAC operation. In section 3, we highlight the major differences between the IEEE 802.14 and MCNS specifications. Section 4 presents the comparative performance evaluation results, and concluding remarks are offered in section 5.

## 2 HFC Overview

In this section we give a brief overview of the HFC system and discuss the main functionality of its PHY and MAC layer protocols. This description is generic and encompasses both the IEEE 802.14 and MCNS specifications.

The HFC system is characterized by a branch and tree topology. At the root of the tree, the headend controls the traffic. The bandwidth is divided into several channels, some dedicated to downstream communication (from the headend to the stations) while others are for upstream transmission (from the stations to the headend). The 802.14 and MCNS standards are focused on the specifications of a PHY as well as a MAC layer protocols in order to implement bi-directional HFC networks.

The PHY specifications define the electrical characteristics of the cable such as the modulation technique, constellation, symbol rates and frequencies used. They also describe several operations performed at the end system physical layer such as scrambling, Forward Error Correction (FEC), ranging and time synchronization.

The MAC protocol that sits above the PHY layer in the protocol stack arbitrates the flow of information between the stations and the headend. Its main responsibility is to ensure that station A is granted permission to send data to the headend without colliding with station B or C or other stations that want to do the same since the upstream channel is a shared medium. A Collision Resolution Protocol (CRP) is invoked in order to resolve collisions resulting from two or more stations transmitting simultaneously.

A MAC Packet Data Unit (PDU) is the basic unit of transfer between the MAC layer at the headend and the station. It consists of a MAC header with or without a data PDU. The same structure is used in both the upstream and downstream directions in order to transmit data and management messages. The upstream channel is divided into discrete time slots called minislots.

A variable number of minislots are grouped to form a MAC layer frame. The headend determines the frame

format by setting the number of data slots (DS) and contention slots (CS) in each frame and sends this information to the stations on the downstream using an *Allocation Message*. Several minislots can be grouped together in order to form a DS that carries a MAC PDU. The DS are explicitly allocated to a specific station by the headend. CS fit into one minislot and are used by the stations to transmit requests for bandwidth. Since more than one station can transmit a request at the same time, CS are prone to collisions. The headend controls the initial access to the CS slots as well as manages the CRP. In order to gain access to the upstream channel a station must follow this multi-step procedure. Upon the arrival of a data packet, the station generates a request and sends it in a CS. In case of a collided CS, the station enters the contention resolution process in order to retransmit its request. On the other hand in case the request is successfully transmitted the station activates its data transmission state machine and exits the contention process. The details of the channel access including the contention process are not the same for the IEEE 802.14 and the MCNS standards and will be described in the next section.

## 3 MAC Protocol: IEEE 802.14 vs MCNS

The MAC protocols described in the IEEE 802.14 and the MCNS specifications are fundamentally different even though they share a number of similar functional requirements. In this section we identify two major differences that may have a direct impact on performance, namely the mapping of higher layer traffic and the upstream channel access policy used which includes the contention slots access and the collision resolution. Most of the other differences we find lie in the messaging formats (request/feedback/grant) and the management and security layers and represent significant conceptual differences between the two standards that make them incompatible. The details of those differences are outside the scope of this paper and can be found in [3] and [4].

### 3.1 Mapping of Higher Layer Traffic

The framing structure of the 802.14 is significantly different from the one adopted in the MCNS specification. Since the 802.14 standard is intended to provide complete support of ATM and in order to minimize the MAC layer overhead, one byte is added to each ATM cell to form a MAC data PDU as shown in Figure 1 where the ATM layer VPI field is used as part of the

14-bit local station ID. In addition, every station must be capable of AAL5 segmentation and reassembly in order to carry IP/LLC traffic. On the other hand the MCNS specifications assume a more friendly IP environment. 6 bytes of MAC header are added to every packet regardless of whether it is an ATM cell or an LLC packet as illustrated in Figure 1. A concatenated mode is also defined for ATM so that multiple of ATM cells can be grouped together with only one MAC header in an attempt to reduce the overhead.

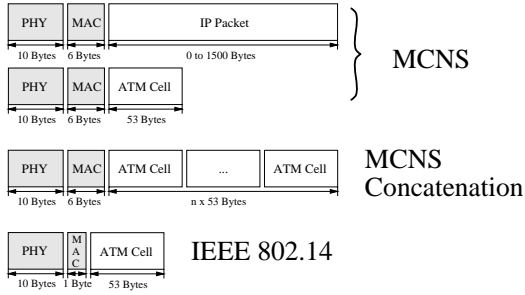


Figure 1: Mapping of Higher Layer

### 3.2 Upstream Contention Access

In the 802.14 standard, the headend tightly controls the initial access to the CS slots as well as manages the CRP by assigning a Request Queue (RQ) number to each CS. Upon receipt of a data packet, the station generates a Request Minislot Data Unit (RMDU). An admission control scheme for newcomer stations is used to provide differentiated initial contention access.

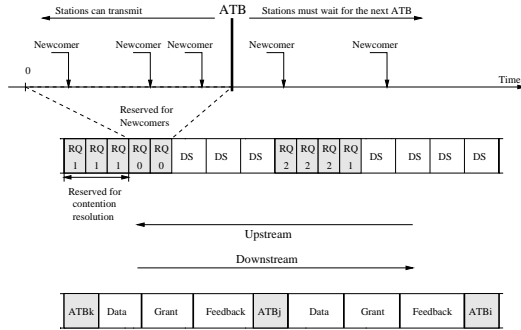


Figure 2: IEEE 802.14 Channel Access

This scheme is based on preassigned priorities and a FIFO service of timestamped requests. For the sake of simplicity in this discussion we assume that all stations have the same priority of access. The headend controls the station's entry by sending an *Admission Time Boundary* (ATB) periodically as illustrated in Figure 2.

Thus only stations with a generated RMDU time less than ATB are eligible to enter the contention process. Once the RMDU is generated, the station waits for a *CS Allocation* message from the headend that reserves a group of CS with  $RQ = 0$  for newcomer transmission. The station randomly selects a CS in that group and transmits its RMDU. Since multiple stations may attempt to send their RMDUs in the same upstream CS a collision may occur. A *Feedback* message is sent to the station after a roundtrip time (which is also equal to a frame length) informing it of the status of the CS slot used. In case of a successful request transmission (*Feedback=Successful*), the station activates its data transmission state machine and exits the contention process. Subsequently a *Data Grant* message will be sent by the headend. In case of a collided CS, the feedback message contains a particular RQ number to be used for collision resolution (*Feedback=RQ*). That is the station needs to retransmit its request in a CS group with that RQ number. The CS groups are usually allocated in the order of decreasing RQ values. For each RQ value the headend assigns a group of CS and an associated splitting value (SPL) that is by default equal to  $3^1$ . A CS within the group is selected randomly in the range  $(0..SPL - 1)$ .

In the MCNS standard, access to the upstream channel is controlled via a backoff window set by the headend (refer to Figure 3). This includes both the initial transmission of a request and any subsequent retransmissions of collided requests. The headend controls the initial access to the CS by setting an initial backoff window, or *Data Backoff Start*. When a station has data to send it sets its internal backoff window equal to *Data Backoff Start* defined in the *Allocation Map* (message equivalent to the *CS Allocation* in the 802.14).

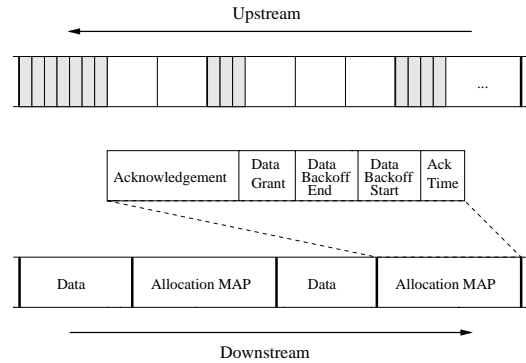


Figure 3: MCNS Channel Access

<sup>1</sup>more details on the ternary-tree based collision resolution used in 802.14 can be found in [2]

The station then randomly selects a number within its backoff window. The random value indicates the number of contention transmit opportunities, which the station must defer before transmitting. After a contention transmission, the station waits for either a *Data Grant* or an *Acknowledgement (Ack)* in a subsequent *Allocation Map* (whichever comes first). Upon receipt of a station's request (in case of a successful transmission), the headend processes it and assigns a DS to the station by sending a *Data Grant* in the *Allocation Map*. The headend may send an *Ack* to the station in case it needs more time to process the request before it sends the *Data Grant*. But since multiple stations may attempt to send their request in the same upstream CS, a collision may occur. However the headend does not need to send an explicit feedback message on the status of each CS as in the 802.14 standard. The station detects the collided slot when it does not find an *Ack* or *Data Grant* for it in the *Allocation MAP*. The station must then increase its backoff window by a factor of two as long as it is less than the maximum backoff window *Data Backoff End* set in the *Allocation Map*. The station randomly selects a number within its new window and repeats the contention process described above. After 16 unsuccessful retries the station discards the MAC PDU.

## 4 Performance Evaluation

We use the *NIST ATM Network Simulator* platform to develop a general simulation model for HFC MAC protocols [5]. We implement both the MCNS and the IEEE 802.14 MAC specifications and use the configuration and system parameters for the HFC network shown in Table 1. Some of our preliminary results on the MCNS and the 802.14 MAC protocols were presented in [7]. All simulations, are run for 30 seconds of simulated time and the first 10% of the data is discarded. Several measurements are used in the performance evaluation, namely, the mean access delay, the coefficient of delay variation, the delay probability density function and the mean cell delay variation (CDV). The access delay is the time it takes a packet to reach the headend from the time the packet is generated at the station. The coefficient of delay variation is the normalized access delay variance, while the delay probability density function gives the distribution of the access delays for a given load. The mean CDV is computed for the CBR traffic in order to measure jitter or the gap between cell arrivals at the destination.

Simulation Parameters	Values
Distance from nearest/furthest station to headend	25/80 km
Downstream data transmission rate	Not considered limiting
Upstream data transmission rates	3 Mbits/sec
Propagation delay	5 $\mu$ s/km for coax and fiber
Length of simulation run	30 seconds
Length of run prior to gathering statistics	10% of simulated time
Physical layer overhead	10 bytes
Guard-band and pre-amble	Duration of 5 bytes
Data slot size	in multiple of 16 bytes
CS size	16 bytes
Frame size	36 mini-slots
Number of CS per Frame	Variable
Roundtrip	1 Frame
Maximum request size	32 DS
Headend processing delay	0 ms
<b>MCNS Parameters</b>	
MAC overhead	6 bytes
ATM transmission size (no concatenation)	53+6+10 bytes
ATM transmission size (w/concatenation)	n*53+6+10 bytes
Variable IP packet transmission size	Frame size+6+10 bytes
CS size	16 bytes
Initial exponential backoff	$2^4=16$ CS
Maximum exponential backoff	$2^8=256$ CS
<b>IEEE 802.14 Parameters</b>	
MAC overhead	1 byte
ATM transmission size	53+1+10 bytes
Variable IP packet transmission size	Segmented into 48 byte cells
Contention Resolution Protocol	Ternary-tree with ATB

Table 1: Simulation Parameters

### 4.1 Experiments

We present the results from four different simulation experiments that stress the difference between the MCNS and the 802.14 systems and evaluate their performance when operating under the same network conditions.

Message Size (bytes)	64	128	256	512	1024	1518
Probability	0.6	0.06	0.04	0.02	0.25	0.03

Table 2: IP Traffic: Message Size Distribution

In **Experiment 1**, we look at the contention resolution algorithms used: binary exponential backoff versus ternary-tree using an ATM cell transfer. Both the framing structure are set to the IEEE 802.14 format, i.e. only 1-byte MAC overhead is added to each ATM cell. The traffic for this test is based on 53-byte ATM cells generated according to a Poisson distribution with

a mean arrival rate of  $\lambda$ , where  $\lambda$  is varied according to the offered load. Short packets or ATM cells are used in this case in order to stress the contention access and demonstrate the effectiveness of the collision resolution algorithm used.

In **Experiment 2**, we study ATM transfer in MCNS and 802.14 networks. The mapping of ATM cells at the MAC layer is set as specified in each standard. We also use the concatenation mode for MCNS which attempts to reduce the MAC overhead by packing multiple ATM cells into one MAC PDU. The traffic distribution is based on the generation of ATM cells according to a Poisson distribution as in Experiment 1.

**Experiment 3** concentrates on IP transfer. In the case of the 802.14, the IP packets are segmented into 48 bytes using AAL5. The traffic is based on bursty sources with a batch Poisson arrival model. The message size distribution is defined according to Table 2. The message interarrival time is exponentially distributed with mean  $T = \frac{1}{\lambda}$ , where  $\lambda$  varies according to the load.

**Experiment 4** focuses on the jitter and end-to-end delay bound sensitive CBR applications in both the MCNS and the 802.14 environments. We measure the cell delay variation available to CBR in a contention based environment. 20 CBR stations send 10% of upstream channel capacity (at a fixed rate of 300 kbits/s). In addition 200 stations generating ATM cells (same as in Experiment 2) are used for the background traffic that is varied from 5% to 75% of the capacity. The effects of the load increase are measured with respect to the Cell Delay Variation (CDV) of the CBR sources. We assume that there is no preemptive scheduling for CBR requests at the headend.

## 4.2 Results

In this section we present the simulation results obtained for the experiments previously described.

**Experiment 1 -** The tree-based collision resolution algorithm and the timestamps used for newcomer access in the 802.14 standard give lower access delays and delay variances than the binary exponential backoff used in the MCNS standard. This is illustrated in Figure 4 where the mean access delay (a) and the coefficient of delay variance (b) are given with respect to the offered load for both standards. As the load is increased from 5% to 84% of the capacity (150-2500 Kbits/s) the difference in access delay between the

802.14 and MCNS becomes significantly larger: 4ms at 1200 Kbits/s, 10ms at 1350Kbits/s and 25ms at 1500 Kbits/s. Note that these differences in delay are only due to the contention resolution factor since the same frame structure is assumed for both experiments. At 1350 Kbits/s the coefficient of delay variance is  $\sim 0.4$  for 802.14 while it is 1.4 for MCNS (Figure 4(b)). This is mainly due to the nature of the binary exponential backoff algorithm used in MCNS that introduces large differences between the minimum and the maximum access delay [1].

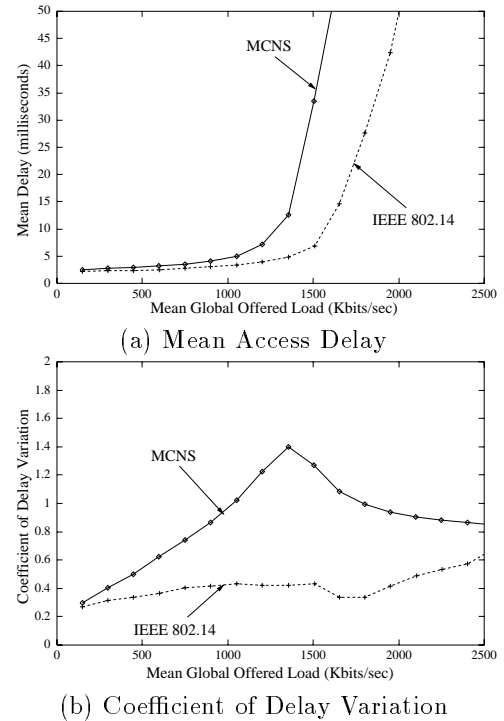
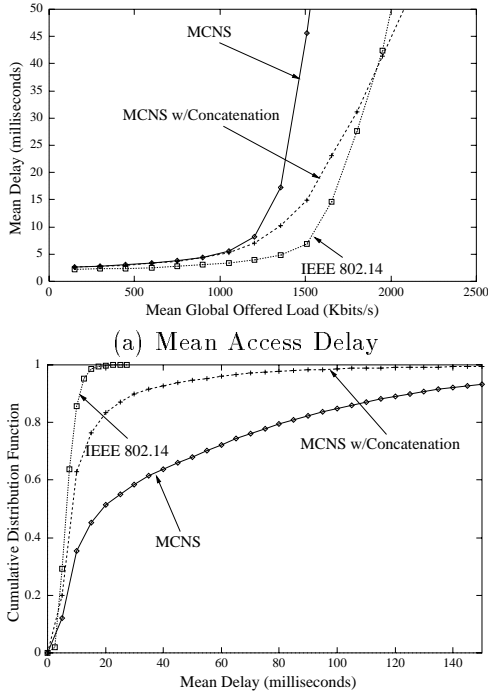


Figure 4: Experiment 1 - Contention Resolution

**Experiment 2 -** In Figure 5 we plot the mean access delay (a) and the delay distribution at 50% of the offered load (b) when using ATM traffic for 802.14 and MCNS. The gap in delay between the MCNS and 802.14 is even wider than in Experiment 1 (Figure 4(a)) due to the differences in ATM mapping between MCNS and 802.14. The MCNS standard adds 16 bytes of overhead to every ATM cell while the 802.14 uses only one byte for the MAC overhead. This results in the access delay curve for MCNS taking off at much lower loads than for 802.14 (around  $\sim 1350$  Kbits/s). At 1500 Kbits/s (50% of capacity) the mean access delay for MCNS is  $\sim 45$ ms while it is  $\sim 6$ ms for 802.14. This represents a significant difference between the two standards with respect to ATM support. The MCNS

concatenated mode for mapping ATM cells reduces the overhead leading to lower delays for higher loads ( $\sim 14\text{ms}$  at  $1500\text{ Kbits/s}$ ). For  $\sim 2000\text{ Kbits/s}$  the delay curves for the MCNS concatenated mode and the 802.14 cross over at  $42\text{ms}$ . In Figure 5 (b) the probability the access delay is less than  $20\text{ms}$  is 1 for 802.14 while it is 0.5 for MCNS and 0.85 for MCNS concatenated mode. The tail of the delay distribution for MCNS is long and the probability distribution does not converge to 1 before  $150\text{ms}$  for both MCNS mapping modes. As observed in Experiment 1, this result is due to the randomness nature of the binary exponential backoff [1]. We note that the distribution of the delay constitute an important measurement in ATM environments where certain traffic types are sensitive to jitter. This measurement was used in [6] to differentiate between the ternary-tree and the p-persistence collision resolution schemes while both were under consideration by the 802.14 group. A decision was made in favor of the ternary-tree because its delay distribution rapidly converged to 1.

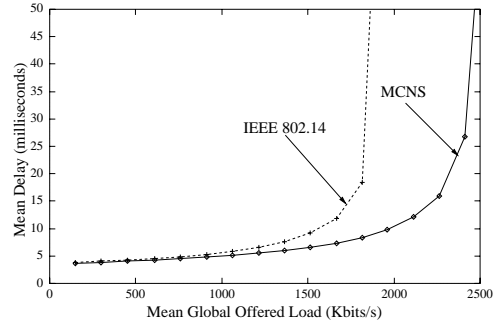


(b) Probability that Mean Access Delay  $< x$  ms  
Offered Load = 50% of Capacity

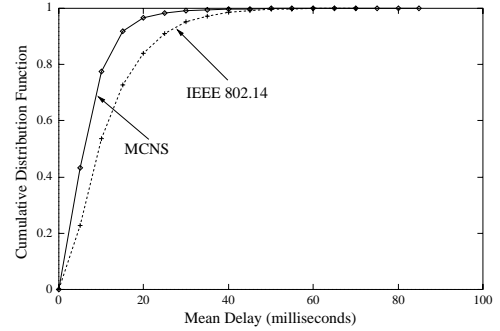
Figure 5: Experiment 2 - ATM Transfer

**Experiment 3 -** When using an IP traffic distribution, the access delay for 802.14 becomes higher than MCNS starting from  $900\text{ Kbits/s}$  (Figure 6(a)). This is the price to pay for AAL5 segmentation of IP packets in

802.14, where for every 48-byte chunk, an overhead (MAC and PHY) of 16 bytes is added. At  $1800\text{ Kbits/s}$  the access delay is  $\sim 17\text{ms}$  for 802.14 while it is  $\sim 7\text{ms}$  for MCNS (Figure 6(a)). From Figure 6(b) the probability the access delay is less than  $20\text{ms}$  is 0.95 for MCNS while it is 0.85 for 802.14 for a load of  $1650\text{ Kbits/s}$ .



(a) Mean Access Delay

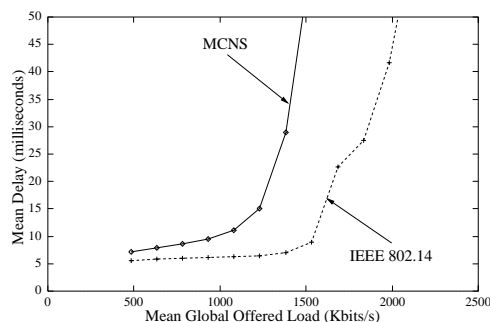


(b) Probability that Mean Access Delay  $< x$  ms  
Offered Load = 55% of Capacity

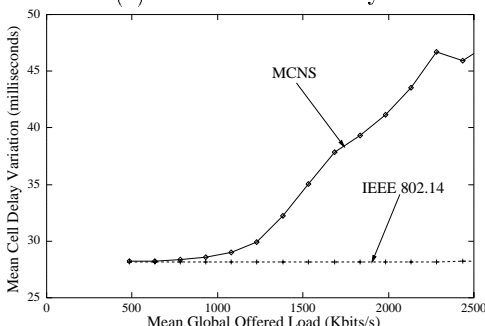
Figure 6: Experiment 3 - IP Transfer

**Experiment 4 -** In Figure 7 we plot the mean access delay (a) and the mean CDV (b) for the 20 stations with CBR traffic. For loads less than  $1500\text{ Kbits/s}$  the delay incurred by the CBR traffic in the 802.14 network is almost constant and is equal to  $\sim 5\text{ ms}$ . This is different in the MCNS network where the mean access delay is increasing from  $5\text{ms}$  to hundreds of milliseconds as the load is increased from  $500$  to  $1500\text{ Kbits/s}$ . These results provide some useful insights on the region of operation of both 802.14 and MCNS cable networks in the presence of delay sensitive traffic. In Figure 7(b) the mean CDV for 802.14 is kept constant at around  $\sim 28\text{ms}$  for loads between  $500$  and  $2500\text{ Kbits/s}$ . This confirms that the 802.14 standard provides good support for CBR type applications such as video and multimedia transfer. However since the mean access delay takes off after 50% of the network capacity ( $1500\text{ Kbits/s}$ ), operating interactive applica-

tions such as videoconferencing and telephony, requires operation in the lower load region (between 0 – 50% of the capacity). The MCNS mean CDV takes off shortly after 1000 kbits/s. At 2250 Kbits/s the CDV for MCNS is around  $\sim 47$ ms which clearly indicates that this standard requires the implementation of special scheduling and management procedures in order to manage CBR traffic.



(a) Mean Access Delay



(b) Mean Cell Delay Variation

Figure 7: Experiment 4: CBR End-to-end CDV

## 5 Concluding Remarks

In this paper we give an overview of the IEEE and MCNS MAC protocols highlighting some of their differences and their impact on performance. We focus on the implementation of the contention access algorithms, the mapping of higher layer traffic and the QOS provision. Given different traffic types we compare the efficiency of both solutions with respect to their access delay and delay variation.

Our simulation results confirm that IEEE 802.14 provides a friendly ATM environment with a good support of QOS while MCNS offers a much more efficient Internet access. The 802.14 solution gives low delays and delay variations for ATM traffic but has lower throughput for IP traffic due to the AAL5 segmentation overhead. The MCNS standard with its rela-

tively simple contention resolution and management algorithms constitute an attractive solution for the transfer of IP traffic. However our results show that more sophisticated mechanisms may be required for better support of QOS.

## References

- [1] D. Bertsekas and R. Gallager. *Data Networks*, 2nd Ed. Prentice Hall, 1992.
- [2] C. Bisdikian, B. McNeil, and R. Norman. msSTART: A random access algorithm for the IEEE 802.14 HFC network. *Computer Communications*, 19(11):876–887, September 1996.
- [3] IEEE 802.14 Working Group. Media Access and Control. IEEE Std 802.14, Draft 3 Revision 1, IEEE 802.14 Working Group, April 1998.
- [4] MCNS Holdings L.P. Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification. SP-RFI-I02-9710008, October 1997.
- [5] N. Golmie, A. Koenig, Y. Saintillan, D. Su. The NIST ATM Network Simulator: Operation and Programming Version 3.0. National Institute of Standards and Technology, Internal Report 5703R1, March 1998.
- [6] N. Golmie, G. Pieris, S. Masson and D. Su. A MAC Protocol for HFC Networks: Design Issues and Performance Evaluation. In *Computer Communications*, volume 20, pages 1042–1050, 1997.
- [7] N. Golmie, M. Corner and D. Su. Choice of a Medium Access Control Protocol for High-Speed Hybrid Fiber/Coax. In *Proceedings of the Ninth IEEE Workshop on Local and Metropolitan Area Networks*, pages 187–192, Banff, Canada, May 1998.